Exploiting Opportunities through Dynamic Coalitions in Robotic Soccer

John Anderson, Ryan Wegner, and Brian Tanner
Autonomous Agents Laboratory
Department of Computer Science
University of Manitoba
Winnipeg, MB, Canada
R3T 2N2

Extended Abstract

While the coordination of multiple agents in static networks or hierarchical relationships has been the subject of much research over a long period of time (e.g. see [1]), the relationships between agents in dynamic multi-agent systems has only comparatively recently been the subject of significant investigation. In dynamic multi-agent systems, the nature of relationships between agents can change, possibly along with the goals of a group. In such systems, the environment surrounding the agents may change, altering the optimality of interactions (e.g. a change in opponent agents; a change in optimal strategy leading to different relationships between agents). While we may have preferences for the way in which agents interact (or even some agents in authority to rely upon), there is no hard-and-fast answer to who should optimally interact with whom and under what conditions. Agents may seek out momentary coalitions with others, or form longer-term bonds that are flexible enough to support some degree of interaction under a wide range of circumstances.

Like multi-agent systems in general, dynamic multi-agent systems cover a spectrum of activity in terms of a number of dimensions: systems may require extremely timely agent decision making, or may have periods where agents can deliberate at length. The position any system occupies within this range influences a number of other factors: the degree to which serious negotiation can occur, for example, obviously depends on the time available for a response. The less such time is available the more agents must rely on knowledge of previous situations or general background knowledge to guide their interactions with others, as opposed to performing active inquiry. In terms of the dynamics of the system, agents may have to rely on those they have been able to trust in the past, rather than actively seeking out new partners when time is of the essence. This is problematic, however, in that in a significantly complex dynamic system, there may be little promise of future results given past performance. Independent of timing, agents may be relatively homogeneous or strongly heterogeneous depending on the nature of the task involved [2]. Depending on the environment, we may have to deal with deception from potential coalition partners, or deception from adversaries influencing our decision on potential collaborators. We may be able to communicate significantly, or communication may have to be limited, clandestine, or indirect.

Our own work in this area arises out of an interest in time-bounded decision making that is representative of the domain of robotic soccer. Robotic soccer has in recent years become a domain of strong interest to the MAS community (see [3,4] to contrast how work in this area has matured) and to AI in general, in that it is representative of a significant number of hard problems in time-bounded reasoning, perception, and agent-interaction, among other areas, and forces interaction between these. Moreover, standard network-based simulators (e.g. [5]) have made this an area that can not only be participated in without a significant investment in robotic equipment, but also one where agents can gain a wide range of experience from global opponents quickly.

Robotic soccer is interesting from the standpoint of dynamic multiagent environments in that even though agents have a high-level role as a member of a team and a shared goal (that is, some permanent structure exists in the relationship between agents), significant dynamic elements exist around this structure. Teamwork and specialization are necessary to do well in soccer: that is, the nature of the task itself leans toward heterogeneity for successful performance [2]. The nature of agent specialization can vary from moment to moment, however: even though we may have positions such as fullback or sweeper, these can be independent of immediate roles in a particular play. Some very successful teams (e.g. [6]) have exploited this by implementing play libraries where roles in particular predefined plays are mapped at run-time according to interests specified by predefined locker room agreements.

While it is certainly possible to have elaborate plays, the selection of these requires strong agreement beforehand simply because of the extremely limited time available in an ongoing soccer game to make selections and map roles through communication. We believe that the extreme time-limited opportunities that characterize robotic soccer make the formation of temporally short coalitions both more interesting from a multi-agent systems standpoint and potentially more successful in the game itself. By a momentary coalition, we mean the explicit recognition that one can be of use to a teammate in an upcoming moment, and the understanding on the part of two or more teammates, verbally or non-verbally, that some opportunity will be exploited between them. An example of such behaviour would be possessing the ball, seeing an open teammate, and receiving a signal from that teammate that they understand the opportunity to achieve a shared goal. A relationship between the two players is momentarily formed over and above any other relationships that currently exist. This relationship may end immediately (I pass the ball), or may continue for some time (I continue up an open channel in the field as a teammate works to keep that channel open). This is as opposed to an elaborate play relying upon understood relationships beforehand and pre-mapped actions as the play unfolds, or on individual behaviours that recognize no relationships with teammates (e.g. kicking the ball forward and someone else happening to be there and moving it further, with no recognized connection on the part of either agent).

Such opportunities are certainly prevalent in human soccer, especially at the midfield position where human agents are expected to get a feel for the overall pace of the game, move the ball up channels, and function to some degree both offensively and defensively. We are working on agents that form momentary coalitions in this manner, and then allow those coalitions to dissipate after the reason for their formation is complete. These agents are being developed in Java using the Robocup soccer server [5] as a test environment, and employ verbal communication through directed and non-directed signals such as *"I'm open"*. This will be extended to non-verbal communication through recognizing a movement pattern in lieu of the server's ability to employ realistic human movements such as waving to an opponent as a signal. We are working largely with two player coalitions, and limited three-player coalitions (a pass while another agent keeps an opponent at a distance, or distracted). This decision is based on the observed behaviour in humans in soccer: in practice, larger spontaneous coalitions in this domain are rare in humans simply because of the problems of recognizing opportunities between several agents in real time and the problems of diverting focus between several other teammates as well as the ball and one's opponents. Once complete, these will serve as a basis for empirically examining the performance of dynamic coalition formation in robotic soccer from the standpoint of parsimony. That is, answer the question of whether such agents can give a performance gain over pure individual behaviours that approaches the use of elaborate large scale plays and the accompanying elaborate representations those require.

This encompasses a number of issues of interest to the community, including:

- The communication and interpretation of intent under time bounds in order to recognize the potential for short-term coalitions;
- Making judgments in real time as to the suitability of other agents for such dynamic coalitions;
- Determining when such a coalition has ended and that its assumed affects on the behaviour of another agent can no longer be relied upon; and
- Developing the greater potential to engage in such coalitions with agents over time as they become more trusted.

The latter is an extremely important issue in its own right that represents a crossover to another area in which we are working, multi-agent learning through peer reinforcement [7]. In this situation, agents are expected to learn a multi-agent task such as the acquisition of behaviours for playing soccer through the reinforcement of peers that are also participating in the same task. However, agents have to gauge the value of reinforcements they receive with their perception of its accuracy, this requires developing trust in the abilities (and in other domains, the motivation) of the agent providing the reinforcements. To this end, learning improves in proportion to the degree that such trust is perceived, and agents must selectively filter reinforcements based on the agent providing it. Within this, we are currently working to examine coalition formation over time and the gradual learning of agents to select as potential coalition partners as well as the settings in which such coalitions are most effective.

We are also working on porting this to scenarios beyond robotic soccer through the more generic player/stage platform [8] provided by the USC robotics laboratory.

Bibliography

[1] Bond, A., and L. Gasser, *Readings in Distributed Artificial Intelligence* (San Francisco: Morgan Kaufmann), 1988. 649 pp.

[2] Balch, Tucker, *Behavioural Diversity in Learning Robot Teams*, Ph.D. Dissertation, Computer Science, Georgia Tech, 1998.

[3] Kitano, H., Tambe, M., and Stone, P., "The RoboCup Synthetic Agent Challenge 97", Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence, Nagoya, Japan, 1997, pp. 24-29.

[4] Stone, P., T. Balch, and G. Kraetzschmar (Eds.), *RoboCup 2000: Robot Soccer World Cup IV* (Heidelberg: Springer-Verlag), 2001. 658 pp.

[5] Noda, I., H. Matsubara, K. Hiraki and I. Frank. "Soccer Server: a Tool for Research on Multi-Agent Systems", *Applied Artificial Intelligence* 12(2-3, 1998):233-250.

[6] Stone, P., and M. Veloso, "Task Decomposition and Dynamic Role Assignment for Real-Time Strategic Teamwork", Proceedings of the 5th International Workshop on Agent Theories, Architectures, and Languages (ATAL-98), pp. 293-308, 1998.

[7] Anderson, J., Tanner, B., and Wegner, R., "Peer Reinforcement in Homogeneous and Heterogeneous Multi-agent Learning", submitted to the 2002 *IASTED International Conference on AI and Soft Computing*, Banff, AB., June, 2002.

[8] Gerkey, B. P., R. T. Vaughan, K. Stoy, A. Howard, G. S. Sukhatme, and M. J. Mataric, "Most Valuable Player: A Robot Device Server for Distributed Control", *Proceedings of the IEE/RSJ International Conference on Intelligent Robots and System*s, Montreal, 2001, pp. 1226-1231.